Bootstrapping Ternary Relation Extractors

Ndapandula Nakashole
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
ndapa@cs.cmu.edu

arXiv:1511.08952v1 [cs.CL] 29 Nov 2015

Abstract

Binary relation extraction methods have been widely studied in recent years. However, few methods have been developed for higher n-ary relation extraction. One limiting factor is the effort required to generate training data. For binary relations, one only has to provide a few dozen pairs of entities per relation, as training data. For ternary relations (n=3), each training instance is a triplet of entities, placing a greater cognitive load on people. For example, many people know that Google acquired Youtube but not the dollar amount or the date of the acquisition and many people know that Hillary Clinton is married to Bill Clinton by not the location or date of their wedding. This makes higher n-nary training data generation a time consuming exercise in searching the Web. We present a resource for training ternary relation extractors. This was generated using a minimally supervised yet effective approach. We present statistics on the size and the quality of the dataset.

1 Introduction

Developing techniques for higher-nary relation extraction is a natural next step after the well-studied case of binary relations (Auer et al., 2007; Suchanek et al., 2007; Bollacker et al., 2008; Carlson et al., 2010; Mitchell et al., 2015). In the literature, prominent binary relation extraction methods are mostly semi-supervised (Suchanek et al., 2009; Carlson et al., 2010) or unsupervised (Mausam et al., 2012; Fader et al., 2011a). Semi-supervised methods tend to have higher precision than unsupervised methods and therefore are commonly used to populate knowledge bases of facts (Nakashole et al., 2011; Mitchell et al., 2015). In such settings, relations of interest are predefined, i.e, company acquisitions or protein-protein interactions. However, in semi-supervised approaches, one needs to provide seed examples for each relation to bootstrap the extractor. This can be expensive, especially if there are many relations of interest. For ternary relations, hand specifying training instances per relation requires even more time since each training instance is a triplet of three entities as opposed to a pair of entities for binary relations. Most people know that Google acquired Youtube but not the dollar amount or the date of the acquisition, and most people know that Hillary Clinton is married to Bill Clinton by not the location or date of their wedding. This makes ternary training data generation a time consuming exercise in searching the Web.

In this paper we present a resource for training ternary extractors. The resource was generated using a minimally supervised yet high precision method. Our method leverages a very common language construction: prepositional phrases (PPs). PPs such as "in X", "at Y", and "for Z" express details about the *where, when,* and *why* of binary relation instances. This makes PPs well-suited to extending binary relations by one more argument, extending them to ternary relations.

Consider the following occurrences of 5-item sequences of the form: N1, V, N2, P, N3; where the Ns are noun phrases, V is a verb and P is a preposition.

(1) ***Mercedes-Benz*** *bought* ***Chrysler*** *for* ***$40 billion***
(2) ***CBS*** *bought* ***WCCO*** *from* ***General Mills***

(3) ***Joe Lieberman*** *endorsed* ***McCain*** *for* ***president***
(4) ***The New Yorker*** *endorsed* ***Obama*** *over* ***Romney***

In a large Web extracted corpus, one sees many 5-item sequences similar to each of the 6 types shown above. The verbs and prepo-

sitions do not change but the arguments N1-3 change. From a large volume of such occurrences, we can learn templates for populating ternary relations. For example, from many occurrences of tuples of type (1), we can generate the template: *<organization>bought <organization>for <dollar_amount>*. We go further by having labels for all three argument placeholders. For this particular template the labels are: *AcquisationEventAcquirer, AcquisationEventAcquired, AcquisationEventAmount* for the N0, N1, N3, argument placeholders respectively. Similarly, from tuples of type (3), we can generate the template: *<person>endorsed <politician>for <political_office>*. And the corresponding argument labels are: *EndorsementEventEndorser, EndorsementEventEndorsed, EndorsementEventOffice*, for the N0, N1, N3 argument placeholders respectively. A triplet of argument labels is considered to be a ternary relation. And matching triplets of entities are considered to be instances of ternary relations.

In summary, the contributions of this paper are twofold:

1) Resource: A resource for boostrapping ternary relation extraction. It contains ternary relations and their instances. It also contains templates used to populate ternary relations. We make the data available for future research. It is also attached as supplementary data to this submission.

2) Data Generation Method: We describe the approach used to generate this resource, which others can replicate to generate similar resources in their domains of interest.

2 Related Work

To supervise binary relation extraction methods, there is an abundance of resources. Knowledge bases (KBs) such as DBPedia (Auer et al., 2007), YAGO (Suchanek et al., 2007), Freebase(Bollacker et al., 2008), and NELL (Mitchell et al., 2015), as well as Open IE tools and resources such as Reverb (Fader et al., 2011b), and Ollie (Mausam et al., 2012) contains many millions of binary relation instances that can be used to distantly train binary relation extraction. However, these knowledge bases are highly impoverished when it comes to ternary relations. In contrast, we provide a resource that can be used directly to train and encourage more research on the higher-nary machine reading and relation extraction. A number of works have studied temporal scoping of facts by adding a time dimension to facts (Wang et al., 2010; Wijaya et al., 2014). While temporal scope generates ternary relations, for example by using reification, this only deals with one type of ternary relation. In contrast, our resource contains 50 ternary relations, spanning 18 high level topics or event types.

ElectionEvent	AwardEvent
HiringEvent	FiringEvent
AcquisitionEvent	WeddingEvent
DivorceEvent	DefeatEvent
MeetingEvent	AttackEvent
ProductLaunchEvent	EarthquakeEvent
MurderEvent	PerformingEvent
SuingEvent	BombingEvent
EndorsementEvent	ShootingEvent

Table 1: Event types (18) or high level topics in our resource. The resource contains 50 ternary relations across these topics.

3 Data Generation

In this section we present our method for generating the ternary relations and their instances.

3.1 Input

As input, our method takes a natural language text corpus and high level topics or even types of interest. Our method automatically learns many different ternary relations relevant for each event type. There are two advantages to specifying event types of interest instead of directly thinking in terms of ternary relations. First, a broad event type can capture many relevant ternary relations that naturally appear in the data. Second, it requires much less human effort to specify one event type than manually specifying a list of all conceivable ternary relations, some of which might not be present in the data.

Table 1 shows the event types covered in our resource. For each event type, we specified at most three *trigger verbs* that indicate a potential mention of the event type. We will later describe how to automatically extend event trigger verbs in an iterative manner. This is done in a similar way to extending seed instances or patterns of binary relations (Suchanek et al., 2009; Nakashole et al.,

2011; Mitchell et al., 2015).

3.2 Candidate Template Generation

Once we have event types defined with their trigger verbs, we can generate ternary relations for each event type. The first step is to generate ternary relation **templates**. An example template is: *<person>endorsed <politician>for <political_office>*. We generate these templates directly from the data. We do this by first parsing the raw corpus, and extracting 5-item sequences of the form: N1, V, N2, P, N3; where the Ns are noun phrases, V is a verb such that V is a trigger verb of one of the events, and P is a preposition. We generate templates from 5-item sequences as follows: we replace every noun phrase N1-N3 by its semantic type. In particular, we do lookups of entity types in two types of semantic hierarchies, WordNet and the NELL type hierarchy (Carlson et al., 2010). We found the two type systems to be complementary: WordNet contains more common nouns whereas NELL contains more proper nouns. Our generated templates can therefore contain a mixture of WordNet types and NELL types. For example, for the *MurderEvent*, the following is a valid template that our approach generated: *<NEL_person>killed <NEL_person>with <WDN_weapon>*. The semantic types in our templates are prefixed by three letters, *NEL* for NELL types, and *WDN* for WordNet types.

We retain, as candidate templates, all the templates of the form: *<N1_type>V <N2_type>P <N3_type>*, whose support size is $>= 3$. That is, the template was generated from three or more 5-item sequences, N1, V, N2, P, N3, with distinct noun arguments (N1-N3).

3.3 Template Filtering

From the candidate templates, a final set of templates is generated. To do this, we manually filter out all templates that do not express useful ternary relations for the topic at hand. Once we have filtered out the templates, for each template, we manually label each template with descriptive labels for the its corresponding noun phrase placeholders. For example, *<NEL_person>killed <NEL_person>with <WDN_weapon>* is labeled with *MurderEventMurderer, MurderEventMurdered, MurderEventInstrument*. Each triple of labels is considered to be a single **ternary relation**. Therefore we have the ternary relation *Murderer_Murdered_Instrument*. Such a ternary relation would has **instances** such as *Bob,Alice,knife*, which indicates that Bob murdered Alice with a knife.

We obtain instances of templates, and hence ternary relations, by retaining the supporting 5-item sequences of each of the accepted templates. Notice that each instance has labeled noun phrases because the instances inherit argument labels from their templates.

It is worth noting that template filtering is the part requiring the most manual supervision. All the other parts are automated. While, the specification of event types and their trigger verbs is also manual, it is quite fast, requiring only up to three trigger verbs per event type.

3.4 Iterative Template Generation

In order to extend the size of the resource, we increase the number of trigger verbs per event type. We do this automatically. First, from the raw data, we extract 5-item sequences of the form $N1, V, N2, AP, N3$. There are two main differences between these sequences and the 5-item sequences we have worked with up to now. First, here V is any verb, not limited to the trigger verbs that were manually specified. Second, we no longer limit the phrase between $N2$ and $N3$ to prepositional phrases, P is now $AP = any\ phrase$. We limit the length of AP to a maximum of three words. We then find 5-item sequences where all three arguments $N1 - 3$ match the arguments of an instance of a template from Section 3.3. Thus, we are using the instances generated so far as distant supervision to discover new templates.

All new (V, AP) pairs forming candidate templates that occur with more than 10 distinct instances of an existing template qualify a new promoted templates. We increase the minimum support required from 3 to 10 to avoid introducing noisy templates. A new template has the same argument role labels as the original template whose instances overlap with the instances of the new template. This process extends the trigger verbs, and is not limited to prepositions for the extraction of the third argument. This also allows the generation of more instances from the newly discovered templates. Notice that the number of

ternary relations remain constant from the initial template generation step where we manually label templates with argument roles.

4 Evaluation

We applied our data generation process to Wikipedia (WKP) and the ClueWeb09 (CWB) corpus. We first generated an initial set of candidate templates from Wikipedia, using the method described in Section 3.2. We did not apply this step to the ClueWeb corpus as it can be noisy. We then manually filtered the generated candidate templates as described in Section 3.3. This is iteration 0. We had a total of 186 templates at iteration 0, and 50 ternary relations across 18 event types. The number of templates increase across iterations. Therefore, in subsequent iterations, we obtain more templates that can be used to populate our 50 ternary relations.

From the initial templates, we generate template instances both from Wikipedia and the ClueWeb corpus. These are triples of entities whose types match the argument types of the template, and they occur with the lexical items that appear in the template. We then use the instances as distant supervision to generate more templates as described in Section 3.4. Again, we only discover new templates from Wikipedia, using only ClueWeb to find instances for the discovered templates. At iteration 1, we discovered an additional 174 templates.

Our method picked up new templates until iteration 3, making a total of 502 templates. Figure 1 shows the cumulative number of templates across iterations. Also shown in Figure 1 is precision of templates. We manually assessed precision at every iteration, we did this by randomly selecting 100 templates discovered at every iteration, or all templates if less than 100 templates were discovered in a given iteration. Precision is 100% at all iterations, except for iteration 2 where it dropped to 95%. This was due to a few cases of semantic drift not being cutoff by thresholds of our methods. This led our method do discover templates with verbs such as "resigned as" in templates associated with *hiring* events, we marked these as wrong.

Figure 3.2 shows the cumulative number of instances picked up at every iteration. We started with 31,161 from WKP and CWB corpora at iteration 0 and ended up with 61,380 instances by

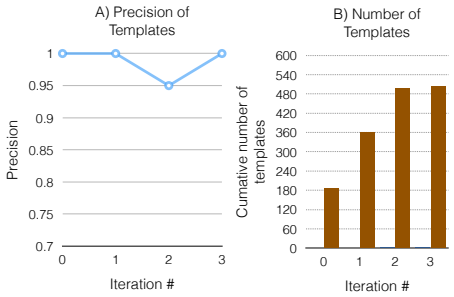

Figure 1: Precision and accumulated number of learned templates, iterations $0-3$

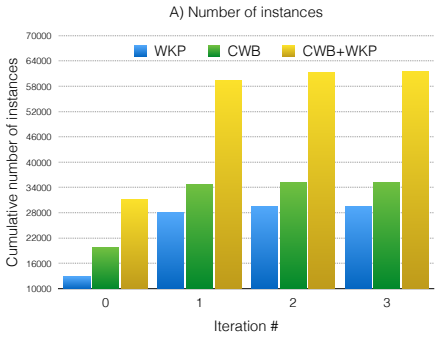

Figure 2: Accumulated number of instances.

the third and final iteration. This number could be increased by 1) allowing discovery of templates from the ClueWeb corpus or other large corpora 2) lowering the high thresholds on the minimum support size of learned templates . Since instances are generated from templates, their precision can be inferred from the precision of the templates. To a small extent, instances can also contains errors stemming from: noun phrase chunking and semantic types; these errors can be fixed by using accurate better chunkers and semantic typing systems.

5 Conclusion

In this paper our goal is to address the bottleneck that has throttled research in higher n-ary relation extraction. To this end, we generated a training data resource for ternary relation extractors. We described a method for learning and populating ternary relations initially only using prepositional phrase based templates. Additionally, our method also learns templates that are not based on prepositions, in an iterative manner. We hope this resource encourages research on ternary relation extraction.

References

[Auer et al.2007] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary G. Ives. 2007. Dbpedia: A nucleus for a web of open data. In *The Semantic Web, 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007, Busan, Korea, November 11-15, 2007.*, pages 722–735.

[Banko et al.2007] Michele Banko, Michael J Cafarella, Stephen Soderland, Matthew Broadhead, and Oren Etzioni. 2007. Open information extraction for the web. In *IJCAI*, volume 7, pages 2670–2676.

[Bollacker et al.2008] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: A collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, pages 1247–1250.

[Carlson et al.2010] Andrew Carlson, Justin Betteridge, Richard C. Wang, Estevam R. Hruschka, Jr., and Tom M. Mitchell. 2010. Coupled semi-supervised learning for information extraction. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, WSDM '10, pages 101–110.

[Del Corro and Gemulla2013] Luciano Del Corro and Rainer Gemulla. 2013. Clausie: Clause-based open information extraction. In *Proceedings of the 22Nd International Conference on World Wide Web*, WWW '13, pages 355–366.

[Fader et al.2011a] Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011a. Identifying relations for open information extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1535–1545. Association for Computational Linguistics.

[Fader et al.2011b] Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011b. Identifying relations for open information extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP '11, pages 1535–1545.

[Kipper et al.2008] Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. 2008. A large-scale classification of english verbs. *Language Resources and Evaluation*, 42(1):21–40.

[Mausam et al.2012] Mausam, Michael Schmitz, Robert Bart, Stephen Soderland, and Oren Etzioni. 2012. Open language learning for information extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, EMNLP-CoNLL '12, pages 523–534, Stroudsburg, PA, USA. Association for Computational Linguistics.

[Mitchell and Lapata2008] Jeff Mitchell and Mirella Lapata. 2008. Vector-based models of semantic composition. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics, ACL*, pages 236–244.

[Mitchell et al.2015] Tom M. Mitchell, William W. Cohen, Estevam R. Hruschka Jr., Partha Pratim Talukdar, Justin Betteridge, Andrew Carlson, Bhavana Dalvi Mishra, Matthew Gardner, Bryan Kisiel, Jayant Krishnamurthy, Ni Lao, Kathryn Mazaitis, Thahir Mohamed, Ndapandula Nakashole, Emmanouil Antonios Platanios, Alan Ritter, Mehdi Samadi, Burr Settles, Richard C. Wang, Derry Tanti Wijaya, Abhinav Gupta, Xinlei Chen, Abulhair Saparov, Malcolm Greaves, and Joel Welling. 2015. Never-ending learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 2302–2310.

[Nakashole and Mitchell2014] Ndapandula Nakashole and Tom M. Mitchell. 2014. Language-aware truth assessment of fact candidates. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 1009–1019.

[Nakashole and Mitchell2015] Ndapandula Nakashole and Tom M. Mitchell. 2015. A knowledge-intensive model for prepositional phrase attachment. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics, (ACL)*, pages 365–375.

[Nakashole and Weikum2012] Ndapandula Nakashole and Gerhard Weikum. 2012. Real-time population of knowledge bases: opportunities and challenges. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction*.

[Nakashole et al.2011] Ndapandula Nakashole, Martin Theobald, and Gerhard Weikum. 2011. Scalable knowledge harvesting with high precision and high recall. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*, WSDM '11, pages 227–236.

[Nakashole et al.2012a] Ndapandula Nakashole, Mauro Sozio, Fabian M. Suchanek, and Martin Theobald. 2012a. Query-time reasoning in uncertain RDF knowledge bases with soft and hard rules. In *Proceedings of the Second International Workshop on Searching and Integrating New Web Data Sources (VLDS at VLDB)*, pages 15–20.

[Nakashole et al.2012b] Ndapandula Nakashole, Gerhard Weikum, and Fabian Suchanek. 2012b. Patty: A taxonomy of relational patterns with semantic types. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, EMNLP-CoNLL '12, pages 1135–1145.

[Nakashole et al.2013] Ndapandula Nakashole, Tomasz Tylenda, and Gerhard Weikum. 2013. Fine-grained semantic typing of emerging entities. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL*, pages 1488–1497.

[Suchanek et al.2007] Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proceedings of the 16th international conference on World Wide Web*, pages 697–706. ACM.

[Suchanek et al.2009] Fabian M. Suchanek, Mauro Sozio, and Gerhard Weikum. 2009. Sofie: A self-organizing framework for information extraction. In *Proceedings of the 18th International Conference on World Wide Web*, WWW '09, pages 631–640.

[Wang et al.2010] Yafang Wang, Mingjie Zhu, Lizhen Qu, Marc Spaniol, and Gerhard Weikum. 2010. Timely yago: harvesting, querying, and visualizing temporal knowledge from wikipedia. In *Proceedings of the 13th International Conference on Extending Database Technology*, pages 697–700. ACM.

[Wijaya et al.2014] Derry Wijaya, Ndapandula Nakashole, and Tom Mitchell. 2014. Ctps: Contextual temporal profiles for time scoping facts via entity state change detection. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.